Set-Based Approach to Design under Uncertainty and Applications to Shaping a Hydrofoil

Johannes O. Royset*
Professor
Operations Research Department, Naval Postgraduate School
Monterey, California 93943
Email: joroyset@nps.edu

Luca Bonfiglio
Postdoctoral researcher
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, Massachusetts 95616

Giuliano Vernengo
Postdoctoral researcher
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, Massachusetts 95616

Stefano Brizzolara
Professor
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, Massachusetts 95616

The paper presents a framework for set-based design under uncertainty and demonstrates its viability through a shape design case of an efficient super-cavitating hydrofoil for ultra-high speed maritime vehicles. Redirecting the effort away from optimal designs to those that safely meet relevant requirements as quantified precisely by superquantile measures of risk (s-risk), the paper addresses situations faced in practice where elimination of poor designs is more important than identifying a best one. The complexity of the design case, involving multiphysics and several design constraints, necessitates the use of surrogate models, which are built using high- and low-fidelity simulations and are here tuned for the first time to reflect s-risk. Accounting for both parameter uncertainty as well as errors in surrogate models, we obtain robust designs that satisfy requirements under a variety of manufacturing and operating conditions.

1 Introduction

It is well known that an optimized design tends to perform poorly under conditions that deviate from those assumed during the optimization. In this paper, we propose a two-pronged approach that addresses this situation and leads to designs that perform robustly under a variety of conditions. The first prong consists of redirecting the effort away from "optimal designs" and instead to focus on identifying sets of candidate designs that meet given requirements. This notion of *set-based design* was pioneered by Toyota and adopted by the U.S. Navy [1]. It responds to most real-world design processes where there are no fixed objective and constraint functions, and therefore no well-defined optimal design problem, but rather an evolving collection of requirements. The goal is then to find designs that can be eliminated from further consideration and, complementarily, find the candidate designs that should be retained. The second prong of our approach is to rigorously account for uncertainty in quantities of interest (QoI) and their models in a manner that captures decision makers' preferences. Here we rely on superquantile measures of risk (s-risk) [2–4] to assess whether QoIs, even in worst-case scenarios, meet the specified requirements.

There is an extensive literature on design under uncertainty; see for example the special issue [5] and the recent articles [4, 6–10]. Overviews of earlier studies are found in [11, 12]. It is beyond the scope of this paper to provide

*Corresponding author

a comprehensive review and comparison. However, we note that much of the literature quantifies performance relative to a threshold by the corresponding failure probability. As laid out in [13], this measure is fundamentally problematic as it is hard to estimate, introduces complexity during optimization, and fails to account for the severity of a failure. The use of s-risk improves on these aspects as described below.

For the first time we consider s-risk in a real design case consisting of a new family of dual-operating mode super-cavitating hydrofoils [14] devised to ensure high efficiency both in super-cavitating and in fully-wet conditions, unlike the conventional geometries. These unconventional hydrofoils are designed to serve as lifting surfaces in a new ultra-high speed Hydrofoil Small Waterplane Area Twin Hull (SWATH) unmanned craft [15] designed to operate as a displacement vessel at low speeds in high sea states and to reach a top speed of more than 100 knots in milder sea states. It is evident that the fully wet operation of the unconventional hydrofoils becomes very attractive during take-off maneuver. In this design case and in fact in most real-world situations the QoIs cannot be computed exactly. Numerical simulations allow us to estimate the QoIs and even that only at high computational costs. We rely on surrogate models, trained by output of the simulations, for approximation of the QoIs across the space of design variables. Deviating from standard surrogate models (see for example [16, 17]), we develop for the first time a risk-tuned surrogate model that is tailored to estimate the s-risk of QoIs. We carry this out in a multi-fidelity context where the training data for the surrogate model comes from a low-fidelity potential flow solver and from a high-fidelity viscous fully turbulent model.

The contribution of the paper is therefore three-fold. First, relying on s-risk, we develop a framework for set-based design under uncertainty in Section 2. Second, we construct a novel surrogate model leveraging multi-fidelity simulations in Section 3. Third, we demonstrate the approach through a complex design case in Section 4, with results in Section 5.

2 Set-Based Design under Uncertainty

In this section, we give a mathematical formulation of the problem of set-based design under uncertainty. The situation is the one usually encountered in practice: the process of design advances through several stages. At an early stage there are many designs and concepts on the table, as well as a number of QoIs and associated requirements. The goal at such a stage is not to identify an "optimal design," but rather a set of candidate designs that satisfies the requirements and thus are worthwhile to consider further. Consequently, the stage *eliminates* designs deemed deficient according to the current requirements. The set of candidate designs advances to the next stage to face a revised collection of QoIs and requirements. The subset of designs meeting these revised requirements becomes the new set of candidate designs and the process is repeated at successively more advanced stages of design. We here provide the mathematical foundations for the problem of identifying the set of candidate design at a stage. Our formulation applies to any stage despite the fact that the stages might involve rather different QoIs and requirements. We do not address the issues of selecting QoIs and requirements as these depend on the design case. Section 4 provides details in the context of shaping a hydrofoil.

2.1 Accounting for Uncertain Parameters

Suppose that at the current stage of the design process there are K QoIs. The numerical value of each QoI is parameterized by a vector $\mathbf{x}$ containing *design variables* and a vector $\mathbf{v}$ consisting of *uncertain parameters.* The design variables specify a design, for example the shape of a hydrofoil as seen below. The goal is to find designs $\mathbf{x}$ that meet the requirements of the current stage. The collection of such $\mathbf{x}$ is the set of candidate designs. The uncertain parameters represent inputs that are beyond the control of the designer and that can affect the performance of the design such as environmental conditions, manufacturing errors, and variability in strength of materials. Their values are not known at the current stage. We let the functions $g_k(\mathbf{x},\mathbf{v})$ give the (true) numerical value of the kth QoI under design $\mathbf{x}$ and parameter $\mathbf{v}$, $k = 1,2,...,K$. Of course, in most practical situations these functions are not available in explicit form, which we address below. Without loss of generality, we assume that low values of the QoIs are preferred to higher values and that the requirement associated with the kth QoI is that $g_k(\mathbf{x},\mathbf{v})$ should not exceed a number r_k. If a practical situation demands high values, we simply replace $g_k(\mathbf{x},\mathbf{v})$ by $-g_k(\mathbf{x},\mathbf{v})$ and r_k by $-r_k$. If both too high and too low values of the kth QoI need to be avoided, then in addition to $g_k(\mathbf{x},\mathbf{v})$ we also consider $-g_k(\mathbf{x},\mathbf{v})$. The former needs not to exceed an upper-bounding requirement and the latter not to exceed the negative of a lower-bounding requirement.

A requirement that a design $\mathbf{x}$ should have $g_k(\mathbf{x},\mathbf{v})$ not exceeding r_k is fundamentally ill-defined because the values of the uncertain parameters $\mathbf{v}$ are not known. We assume that the uncertainty in $\mathbf{v}$ can be quantified by a known probability distribution. Typically, such distributions are estimated based on historical data about the parameters, but we omit a discussion of this subject here; see for example [18]. When viewing the uncertain parameters as a random vector with a given distribution, we denote it by $\mathbf{V}$. A realization of $\mathbf{V}$ is still denoted by $\mathbf{v}$. Now $g_k(\mathbf{x},\mathbf{V})$ is a random variable for each $\mathbf{x}$, which allows us to bring in s-risk.

S-risk [13] reduces a random variable to a number, which somehow represents the random variable and can be used in comparison with requirements. It was originally proposed for finance applications under the name conditional value-at-risk [2, 3]. Specifically, for *risk-averseness parameter* $\alpha \in [0,1]$, the s-risk of $g_k(\mathbf{x},\mathbf{V})$ at level α, denoted by $R_\alpha(g_k(\mathbf{x},\mathbf{V}))$, is understood as the

$$\begin{aligned}&\text{average value of } g_k(\mathbf{x},\mathbf{V}) \text{ over the worst}\\ &(1-\alpha)100\% \text{ outcomes of } g_k(\mathbf{x},\mathbf{V}).\end{aligned} \tag{1}$$

This expression can be taken as the definition if $g_k(\mathbf{x},\mathbf{V})$ has a continuous distribution function. If the risk-averseness

parameter $\alpha = 0$, then $R_\alpha(g_k(\mathbf{x},\mathbf{V}))$ is simply the mean of $g_k(\mathbf{x},\mathbf{V})$. If $\alpha = 1$, then $R_\alpha(g_k(\mathbf{x},\mathbf{V}))$ is the largest possible value of $g_k(\mathbf{x},\mathbf{V})$. Of course, a value of α between these two extremes provides a middle ground. If the random variable $g_k(\mathbf{x},\mathbf{V})$ is normal with mean μ and standard deviation σ, then $R_\alpha(g_k(\mathbf{x},\mathbf{V})) = \mu + \sigma\varphi(\Phi^{-1}(\alpha))/(1-\alpha)$, where φ is the standard normal probability density function and Φ is the standard normal cumulative distribution function. We refer to [4] for general methods for computing s-risk. Regardless of the situation, s-risk is fundamentally appealing as it has a simple interpretation, accounts for the tails of probability distributions, and introduces in some sense no significant complexity compare to a problem without uncertainty. Specifically, if $g_k(\mathbf{x},\mathbf{v})$ is convex in $\mathbf{x}$, then $R_\alpha(g_k(\mathbf{x},\mathbf{V}))$ will also be convex in $\mathbf{x}$. This implies that finding $\mathbf{x}$ with $R_\alpha(g_k(\mathbf{x},\mathbf{V})) \le r_k$ is in some sense no harder than finding an $\mathbf{x}$ with $g_k(\mathbf{x},\mathbf{v}) \le r_k$ for a fixed $\mathbf{v}$. In contrast, finding an $\mathbf{x}$ such that the (failure) probability of $g_k(\mathbf{x},\mathbf{V})$ exceeding r_k is fundamentally hard. This probability also fails to account for the magnitude of any exceedance of r_k; see [4, 13].

We are now in a position to formalize the requirement that the value of the kth QoI should not exceed r_k. Specifically, we say that $g_k(\mathbf{x},\mathbf{V})$ is *safely* $\le r_k$ if $R_\alpha(g_k(\mathbf{x},\mathbf{V})) \le r_k$. In view of Eqn. (1), it is clear that this interpretation of the requirement implies that even on average over a set of worst outcomes, the value of the QoI will not exceed r_k. The specific definition of "worst outcomes" depends on α. We can also interpret the requirement probabilistically. If we let the event that $g_k(\mathbf{x},\mathbf{V})$ exceeds r_k be define as *failure*, then $R_\alpha(g_k(\mathbf{x},\mathbf{V})) \le r_k$ is equivalent to having the *buffered failure probability* $\le 1-\alpha$ as we see from [13]. That reference also establishes that $R_\alpha(g_k(\mathbf{x},\mathbf{V})) \le r_k$ implies that the failure probability is $\le 1-\alpha$. Letting $\alpha_k \in [0,1]$ be the risk-averseness parameter associated with the kth QoI, the above discussion leads to the following requirements on the design $\mathbf{x}$:

$$R_{\alpha_k}(g_k(\mathbf{x},\mathbf{V})) \le r_k \text{ for all } k = 1,2,...,K.$$

This formulation accounts for the uncertainty in the parameters $\mathbf{V}$. We build on this as we proceed next to also account for uncertainty in surrogate models.

2.2 Accounting for Uncertainty in Surrogate Models

When the probability distribution of $g_k(\mathbf{x},\mathbf{V})$ is known, it is straightforward to compute $R_{\alpha_k}(g_k(\mathbf{x},\mathbf{V}))$; see [4]. However, this is rarely the case as real-world situations involve most often complex functions $g_k(\mathbf{x},\mathbf{v})$. In particular, the probability distribution of $g_k(\mathbf{x},\mathbf{V})$ is usually unknown even if that of $\mathbf{V}$ is given. To overcome this difficulty we utilize *surrogate models*.

S-risk of $g_k(\mathbf{x},\mathbf{V})$ is a function of the design $\mathbf{x}$, which we denote by $s_k(\mathbf{x})$, i.e., $s_k(\mathbf{x}) = R_{\alpha_k}(g_k(\mathbf{x},\mathbf{V}))$. When this function is inaccessible, as is usually the case, we develop a surrogate model of $s_k(\mathbf{x})$ denoted by $S_k(\mathbf{x})$. Thus, in some sense $s_k(\mathbf{x}) \approx S_k(\mathbf{x})$. We permit that S_k is a random field over the space of design variables, with the randomness arising from the uncertainty associated with the accuracy of the surrogate model. Examples of surrogate models in the form of random fields of this kind are those based on kriging and co-kriging; see for example [16, 17]. In the next section we provide another class of surrogate models that are also of this form.

The original requirement $s_k(\mathbf{x}) \le r_k$, which is usually impossible to evaluate, is now replaced by one involving the surrogate model. In a situation analogous to that above we find that requiring $S_k(\mathbf{x}) \le r_k$ is ill-posed as $S_k(\mathbf{x})$ is a random variable for every $\mathbf{x}$. Thus, again, we turn to s-risk. For risk-averseness parameters $\beta_k \in [0,1]$, $k = 1,2,...,K$, we adopt the requirements

$$R_{\beta_k}(S_k(\mathbf{x})) \le r_k \text{ for all } k = 1,2,...,K.$$

This is a well-posed requirement with the interpretation that

$$R_{\beta_k}(S_k(\mathbf{x})) = \text{the average value of } S_k(\mathbf{x}) \text{ over the worst } (1-\beta_k)100\% \text{ outcomes of } S_k(\mathbf{x}).$$

We stress that here the average is over possible values of an uncertain surrogate model. These values are unrelated to the uncertain parameters $\mathbf{V}$. Since the surrogate model is imperfect, $S_k(\mathbf{x})$ is random and we account for this fact when making the comparison with r_k.

We are then in a position to define the *set of candidate designs* at the current stage:

$$\text{designs } \mathbf{x} \text{ with } R_{\beta_k}(S_k(\mathbf{x})) \le r_k \text{ for all } k = 1,2,...,K. \quad (2)$$

Designs that meet these requirements are *safeguarded* against poor performance, accounting for both parametric uncertainty in $\mathbf{V}$ (using risk-averseness parameters α_k) and surrogate model imprecision (with risk-averseness parameters β_k). Designs that do *not* meet these requirements are eliminated from the design process. The set of candidate designs are carried forward to the next design stage, where they will faced revised requirements for possibly new QoIs. The process is then repeated, possibly with revised risk-averseness parameters.

3 Risk-Tuned Surrogate Models

The previous section permits any surrogate model $S_k(\mathbf{x})$. Here for the first time we develop a surrogate model that is tuned to estimating s-risk; its theoretical foundation is given in [19]; see also [20]. For simplicity, we omit the subscript k as the process will be identical for each QoI.

For a QoI given by $g(\mathbf{x},\mathbf{v})$, typically from a high-fidelity simulation, we leverage a low-fidelity approximation $h(\mathbf{x},\mathbf{v})$. There is no difficulty with including many levels of fidelity, but we omit the details. We are given

a data set $\{(g_i, h_i, \mathbf{x}_i, \mathbf{v}_i)\}_{i=1}^N$ consisting of designs $\mathbf{x}_i$, parameter values $\mathbf{v}_i$, and the corresponding simulation output $g_i = g(\mathbf{x}_i, \mathbf{v}_i)$ and $h_i = h(\mathbf{x}_i, \mathbf{v}_i)$. We adopt a model such as $c_0 + c_1 h_i + c_2 h_i^2 + \mathbf{c}^\top \mathbf{x}_i$ to approximate g_i, where c_0, c_1, and c_2 are scalar coefficients and $\mathbf{c}$ is a vector of coefficients to be determined by a particular type of error minimization. Of course, models with more or less terms can easily be considered. The generation of the data set requires high- and low-fidelity simulations of the relevant QoI.

The essence of the surrogate model is the method by which the coefficients in the model are fitted. Specifically, let $\alpha \in (0,1)$ be the risk-averseness parameter of interest. Then, the proposed method stipulates that the coefficients should be determined through α-quantile regression, i.e., find coefficients c_0, c_1, c_2, and $\mathbf{c}$ that minimize

$$\frac{1}{(1-\alpha)N}\left(\sum_{i=1}^{N} \max\{0, e_i\}\right) + \frac{1}{N}\sum_{i=1}^{N} e_i,$$

where the error $e_i = g_i - (c_0 + c_1 h_i + c_2 h_i^2 + \mathbf{c}^\top \mathbf{x}_i)$. We discard the optimal c_0, but name the other optimized coefficients $\hat{c}_1$, $\hat{c}_2$, and $\hat{\mathbf{c}}$. Finally, we set $\hat{c}_0 =$ the s-risk of $\{g_i - (\hat{c}_1 h_i + \hat{c}_2 h_i^2 + \hat{\mathbf{c}}^\top \mathbf{x}_i)\}_{i=1}^N$. Corollary 4.2 in [19] states that

$$R_\alpha(\{g_i\}_{i=1}^N) \le R_\alpha(\{\hat{c}_0 + \hat{c}_1 h_i + \hat{c}_2 h_i^2 + \hat{\mathbf{c}}^\top \mathbf{x}_i\}_{i=1}^N).$$

Consequently, over the given data set the model $\hat{c}_0 + \hat{c}_1 h(\mathbf{x}, \mathbf{v}) + \hat{c}_2 h^2(\mathbf{x}, \mathbf{v}) + \hat{\mathbf{c}}^\top \mathbf{x}$ leads to an upper approximation of the s-risk of the QoI. This motivates the use of the model as an approximation of the s-risk of the QoI, i.e.,

$$R_\alpha(g(\mathbf{x}, \mathbf{V})) \approx R_\alpha(\hat{c}_0 + \hat{c}_1 h(\mathbf{x}, \mathbf{V}) + \hat{c}_2 h^2(\mathbf{x}, \mathbf{V}) + \hat{\mathbf{c}}^\top \mathbf{x}). \quad (3)$$

There are two reasons for why this approximation might fall short of an actual inequality. First, fitting of the coefficients of the surrogate model might be carried out using a data set of size N and not the true distribution of the random vector $\mathbf{V}$. Second, Eqn. (3) makes comparisons at individual designs $\mathbf{x}$, which might generally be different than considering "averages" across a variety of designs as relied on while fitting the coefficients.

We note that it is central that the risk-averseness parameter α used in the regression matches that of interest so that the coefficients in the model are "tuned" appropriately. In the case of $\alpha = 0$, we simply determine the coefficients by means of standard least-squares regression, and for $\alpha = 1$ see [19].

It would also be prudent to account for the possible inaccuracy in the approximation in Eqn. (3). We achieve this by replicating the above estimation on different data sets such that not only one but J sets of coefficients are determined. For $j = 1, ..., J,$, let

$$S^{(j)}(\mathbf{x}) = R_\alpha\left(\hat{c}_0^{(j)} + \hat{c}_1^{(j)} h(\mathbf{x}, \mathbf{V}) + \hat{c}_2^{(j)} h^2(\mathbf{x}, \mathbf{V}) + (\hat{\mathbf{c}}^{(j)})^\top \mathbf{x}\right)$$

be the resulting models. Finally, we define the surrogate model as $S(\mathbf{x}) = \{S^{(j)}(\mathbf{x})\}_{j=1}^J$, which can be viewed as a random field over the space of design variables, with J possible and equally likely realizations. As stipulated in Eqn. (2), a risk-averseness parameter $\beta \in [0,1]$ and a corresponding s-risk lead to the condition $R_\beta(S(\mathbf{x})) \le r$ for designs to satisfy, where r is a requirement.

In summary, we see that through Eqn. (3), the various $S^{(j)}(\mathbf{x})$ approximate the s-risk of the quantity of interest $R_\alpha(g(\mathbf{x}, \mathbf{V}))$. However, to account for the fact that these approximations are uncertain, we consider the average of the worst-case values of $S^{(j)}(\mathbf{x})$ in some sense that depends on β. That average is compared with the given requirement. For example, if $\beta = 1$, then $S(\mathbf{x}) = \max_{j=1,...,J} S^{(j)}(\mathbf{x})$ and the worst-case model is used in comparison with r. For lower β, an average of the higher $S^{(j)}(\mathbf{x})$ would be used.

4 Hydrofoil Design

As a practical application of the proposed method, we selected the design, by numerical simulation, of a new surface-piercing super-cavitating hydrofoil. The performances achievable with a CFD based optimization offer hydrofoil crafts new possibilities of reaching higher speeds at sea avoiding the hysteresis phenomena typically affecting conventional super-cavitating profiles. The innovation of this hydrofoil lies in the shape of the 2D sections, which are designed to ensure particular hydrodynamic properties. Traditional super-cavitating hydrofoils are designed to operate at very high speeds only where flow conditions ensure the presence of a stable vapor cavity enveloping the entire suction surface and closing many chords aft a blunt trailing edge. The shape of the unconventional hydrofoil presented in [14, 21] features a pointed leading edge ensuring cavity detachment at the operating angle of attack and cavitation index in design conditions and a sharp edge on the pressure side triggering base cavitation when working at zero angle of attack. The blunt trailing edge, typical of conventional super-cavitating profiles, is instead tapered in a tail which is designed to be enclosed in the supercavity at high speed while producing a good pressure recovery and higher lift at lower speeds when cavitation disappears. This guarantees the performance of the hydrofoil also in sub-cavitating conditions. The tail is functionally separated by the forward body of the profile though a sharp corner (face cavitator) on the face in order to trigger base cavitation at intermediate speeds or at high speeds and lower angles of attack. Control points of composite B-Spline curves provide a fully parametric description of the shape.

A benchmark hydrofoil has been designed by means of a differential evolution optimization algorithm (DE) in [22], where CFD simulations were performed using a low-order Boundary Element Method (BEM) to solve for the steady potential flow with cavitation. Penalty terms were added to the objective function when design variables and/or QoIs fell outside of given bounds. No uncertainty was considered in the development of the benchmark design. Figure 1 presents a time-averaged snapshot of the flow around the benchmark

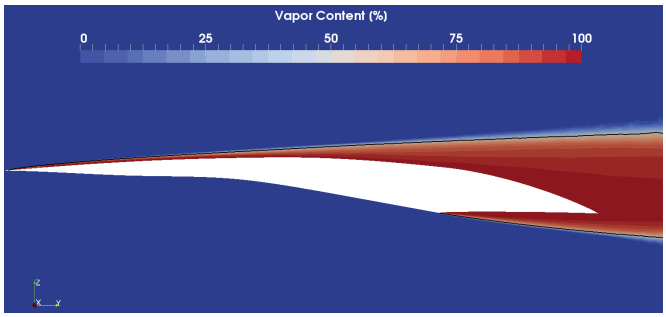

Fig. 1. Time-averaged vapor content contours for the benchmark hydrofoil resulting from an optimization without accounting for uncertainty. Flow follows positive y direction (left to right). Red colors indicate stable cavity enveloping the suction surface and detaching from the face cavitator (sharp edge in the lower hydrofoil surface) with Reynolds number $4.58 \cdot 10^7$, angle of attack 6 degrees, and cavitation index 0.05.

design. The cavity shape can be inferred from the color scale indicating vapor content. The optimization leads to a thin cavity enveloping the hydrofoil surfaces showing the drawback of the design: it is highly sensitive to changes in operating conditions. In this case the existence of a super-cavity relies on the presence of a vapor regime over the hydrofoil suction surface. It is evident that manufacturing errors or variable conditions (especially in terms of angle of attack) could significantly affect the flow regime and hydrofoil performances for a design with a thin cavity. This recognition leads us to considering the set-based design framework under uncertainty described above. The remainder of the paper demonstrates that this framework produces robust designs that perform well under a range conditions.

4.1 Description and Parameters

We provide here a detailed description of the hydrofoil design problem. The design vector $\mathbf{x}$ consists of 15 variables representing coordinates of the control polygon of the composite B-spline curves describing the shape of a hydrofoil. The configuration of the control points has been chosen in such a way that all desired shape variations are allowed both on the suction and pressure side. Figure 2 gives a schematic representation of the 15 variables together with a resulting hydrofoil shape. The pressure side is described by five control points, one of which allows to change the vertical position of the face cavitator, i.e., the sharp edge at the beginning of the hydrofoil tail. Variations of the suction side are performed by five control points; three of them move together vertically following the coordinate ($\bar{x}_{15}$). The leading and trailing edge position as well as the control points of the tail pressure side are kept fixed. The control points on the hydrofoil face regulate lift and drag performance while those on the back suction side are needed to control the cavity thickness and maintain a sufficient strength (inertia modulus).

We consider uncertainty related to manufacturing errors. Specifically, a random vector $\mathbf{V} = (V_1, ..., V_{15})$ describes relative displacement values to the control point coordinates.

Table 1. Benchmark design and bounds of design variables.

Variable	Benchmark	$x_i^{\min}$	$x_i^{\max}$
x_1	$5.556 \cdot 10^{-2}$	$5.556 \cdot 10^{-2}$	$1.667 \cdot 10^{-1}$
x_2	$2.227 \cdot 10^{-1}$	$2.222 \cdot 10^{-1}$	$3.333 \cdot 10^{-1}$
x_3	$3.900 \cdot 10^{-1}$	$3.889 \cdot 10^{-1}$	$5.000 \cdot 10^{-1}$
x_4	$5.576 \cdot 10^{-1}$	$5.556 \cdot 10^{-1}$	$6.667 \cdot 10^{-1}$
x_5	$5.556 \cdot 10^{-3}$	0	$5.556 \cdot 10^{-3}$
x_6	$1.111 \cdot 10^{-2}$	$1.111 \cdot 10^{-3}$	$1.111 \cdot 10^{-2}$
x_7	$3.456 \cdot 10^{-2}$	$1.333 \cdot 10^{-2}$	$3.756 \cdot 10^{-2}$
x_8	$1.833 \cdot 10^{-2}$	$5.556 \cdot 10^{-3}$	$2.778 \cdot 10^{-2}$
x_9	$5.556 \cdot 10^{-3}$	$-1.111 \cdot 10^{-3}$	$5.556 \cdot 10^{-3}$
x_{10}	$1.979 \cdot 10^{-1}$	$1.977 \cdot 10^{-1}$	$2.417 \cdot 10^{-1}$
x_{11}	$4.208 \cdot 10^{-1}$	$4.000 \cdot 10^{-1}$	$4.889 \cdot 10^{-1}$
x_{12}	$5.604 \cdot 10^{-1}$	$5.500 \cdot 10^{-1}$	$6.722 \cdot 10^{-1}$
x_{13}	$3.833 \cdot 10^{-2}$	$3.200 \cdot 10^{-2}$	$3.911 \cdot 10^{-2}$
x_{14}	$6.711 \cdot 10^{-2}$	$5.600 \cdot 10^{-2}$	$6.844 \cdot 10^{-2}$
x_{15}	$8.267 \cdot 10^{-2}$	$7.044 \cdot 10^{-2}$	$8.267 \cdot 10^{-2}$

That is, the actual shape manufactured from a (nominal) design $\mathbf{x} = (x_1, ..., x_{15})$ is assumed to be $\bar{x}_i = x_i(1 + V_i)$ for all $i = 1, ..., 15$. We refer to the vector $\bar{\mathbf{x}} = (\bar{x}_1, ..., \bar{x}_{15})$ as the (random) shape of the hydrofoil under design $\mathbf{x}$. The designer selects the design $\mathbf{x}$ from an allowable range given by the bounds $x_i^{\min} \leq x_i \leq x_i^{\max}$ for all $i = 1, ..., 15$; see Table 1, which also gives the control point values of the benchmark design in Figure 1. The values of the design variables are relative to the hydrofoil chord.

The distribution of $\mathbf{V}$ is discrete with 898 equally likely realizations. These realizations are generated by sampling 5000 15-dimensional vectors from a uniform distribution on $[-0.06, 0.06]^{15}$ and then eliminating realizations that are deemed unrealistic. Thus, "noise" up to $\pm 6\%$ is added to x_i. The framework for set-based design under uncertainty is indifferent to the distribution of $\mathbf{V}$. Here, we select a discrete distribution to avoid errors due to sampling and other approximations, and more easily highlight the features of the framework. Using chord length $c = 0.66$m, we define five ($K = 5$) QoIs: the first being the profile inertia w that is required to exceed the minimum value $w^{\min} = 8.1 \cdot 10^{-6}$ m^3. The second QoI is the profile thickness t_P^2 at 2% of the chord that needs to exceed the minimum value $t_P^{\min}/c = 0.2\%$. The third QoI defines hydrofoil performances in terms of drag-over-lift ratio C_D/C_L that needs to meet a requirement of 0.1, where $C_D = D/(0.5\rho c U^2)$ and $C_L = L/(0.5\rho c U^2)$ normalize drag D and lift L under operating speed $U = 61.667$ m/s and water density $\rho = 997.3$ kg/m^3. The fourth QoI is lift C_L having a requirement of 0.2625 at the design angle of attack $\alpha_{des} = 6$ degrees and at the design cavitation number $\sigma_{des} = 0.05$. The fifth QoI is that the negative lift $-C_L$

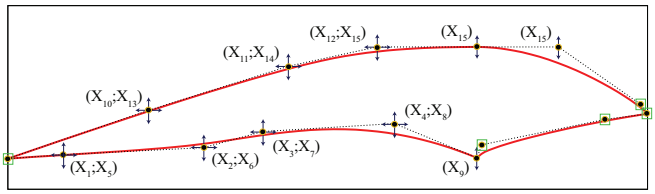

Fig. 2. Definition of design variables. Horizontal and vertical arrows indicate the allowed movement of each control point, and correspond to design variables $\mathbf{x} = (x_1, ..., x_{15})$. Some additional control points are fixed and illustrated by green squares. Three control points on the upper suction side move together according to x_{15}.

must meet the requirement -0.23. The last two QoIs implements the condition that $C_L \in [0.23, 0.2625]$. The risk-averseness parameters are $\alpha_1 = \alpha_2 = 0.7$ (for the geometric properties), $\alpha_3 = 0.9$ (for drag-over-lift), and $\alpha_4 = \alpha_5 = 0$ (for lift). Consequently, we seek designs that on average in the worst 30% of the realizations of $\mathbf{V}$ has $-w \leq -w^{\min}$ and $-t_P^2 \leq -t_P^{\min}$. Moreover, on average in the worst 10% of the realizations of $\mathbf{V}$ we have $C_D/C_L \leq 0.1$ and the average of C_L is between 0.23 and 0.2625. Although suppressed by the notation, w, t_P^2, C_D, and C_L depend on the shape of the profile and therefore on both the design $\mathbf{x}$ and the random vector $\mathbf{V}$.

The first two QoIs will not need surrogate models as there are explicit formula for inertia and profile thickness. For drag-over-lift and lift we develop surrogate models of the kind given in Section 3 using $\beta_3 = 0.95$ for drag-over-lift and $\beta_4 = \beta_5 = 0.$ for lift. The number of replications of surrogate models is set to $J = 9$, each fitted using roughly a ninth of the data as described below. The choices of risk-averseness parameter β_3 then implies that the value of the worst surrogate model over these 9 is used in comparison with the requirement of 0.1. For the lift, the average is employed.

4.2 Analysis Models and Simulations

We have to our disposal a high-fidelity solver for an unsteady viscous method for the solution of Reynolds Averaged Navier-Stokes equations (URANSE) as well as a low-fidelity potential flow solver that both return for a given shape $\bar{\mathbf{x}} = (\bar{x}_1, ..., \bar{x}_{15})$ the lift C_L and the drag C_D at the operating conditions. The URANSE and potential flow output are referred to as high- and low-fidelity output, respectively. In the present context, each high-fidelity solution takes about 6 hours on 4 cores, while each low fidelity run requires 5 seconds on a single core. Figure 1 displays high-fidelity output.

The potential flow solver is based on a BEM in which a Laplace equation is solved for a perturbation potential and a Bernoulli equation is solved instead of continuity and momentum equations for a steady non-viscous flow. Friction forces on the wetted surface of the hydrofoil are included using an empirical local friction coefficient. A cavitation model solving for the cavity thickness on the hydrofoil suction surface is included by imposing an additional set of kinematic boundary conditions that are satisfied through additional singularities distributed on the hydrofoil suction surface and wake panels. This simplified model is able to predict mid-chord cavity closure. Since face cavitation does not represent a feasible engineering design, bubble dynamic has not been modeled for cavities developing on the hydrofoil pressure surface; see [22] for a detail description of the potential flow solver in this setting.

The fully turbulent solution of URANSE is achieved for a multi-phase flow consisting of a fluid mixture of liquid and vapor. The problem is numerically approached using a finite volume technique with collocated arrangement of variables where PDEs are solved for the fluid mixture through a scalar quantity γ indicating the relative content of vapor with respect to liquid within each cell. This surface capturing method developed by [23] is known as Volume of Fluid technique and it is particularly suitable for the solution of complex fluid interfaces as the one characterizing unstable cavitating flows. An additional transport equation for γ, representing the vapor content, is solved together with RANS equations. This transport equation is solved once the source term indicating the specific net mass transferred in the cavitation process is known. Condensation and vaporization regulating the above mentioned mass transfer are found through the cavitation model formulated by [24]. In this simplified method the uncompressible assumption applied to the liquid phase is extended to the vapor phase, hence no additional thermodynamic equations are required to model the phase change which relies only on the local pressure within each control volume. The super-cavitating flow around the hydrofoils is solved using a hybrid structured-unstructured grid consisting in three different regions: a structured grid is used close to the hydrofoil surface to accurately predict the strong gradients characterizing boundary layer flows, two nested unstructured grids of tetrahedral elements are used to increase the resolution in the region of cavity development (hydrofoil wake) while keeping a coarse discretization in the far field region; see [25] for details.

For the shapes resulting from the benchmark design with the 898 realizations of $\mathbf{V}$ we run the URANSE solver and obtain C_D and C_L for all 898 resulting shapes. The same input to the potential flow method results in some estimates of C_D and C_L, but also approximately 50% computational failures, due to the predicted presence of face midcord cavitation that is not presently resolved by the method. Consequently, instead of using the actual low-fidelity simulations for each of the 898 shapes, we compute for each shape a weighted average over all the successful runs, with weights determined by the distance from the current shape to the shape of those runs. Throughout, we deal exclusively with this modified low-fidelity data set. Figure 3 illustrates the high- and low-fidelity drag-over-lift values across the 898 shapes. These runs generate the data from which the surrogate models are fitted according to the procedure in Section 3. Since $\alpha_3 = 0.9$, the surrogate model for drag-over-lift is achieved by carrying out quantile regression with $\alpha = \alpha_3$. The data set of 898 drag-over-lift values from the high- and low-fidelity simulations is divided into $J = 9$ roughly equally large groups, on which quantile regression is car-

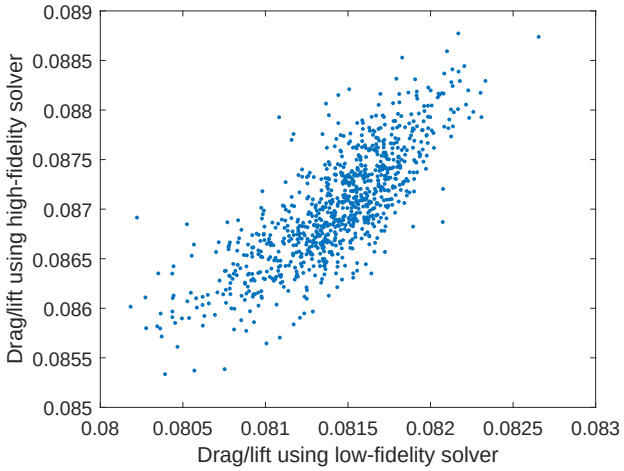

Fig. 3. Low- and high-fidelity estimates of drag-over-lift. Values in meter.

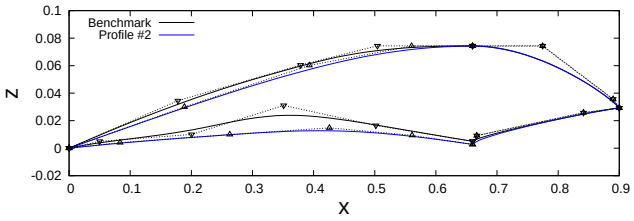

Fig. 4. Design 2 (blue curve, triangles for control points) compared with benchmark profile (black curve, squares for control points).

ried out and a surrogate model $S_{D/L}^{(j)}(\mathbf{x})$ is computed for each $j = 1,...,9$. Since we would like to account for the variability in the models, we adopt the whole ensemble as the surrogate model for drag-over-lift, i.e., $S_{D/L}(\mathbf{x}) = \{S_{D/L}^{(j)}(\mathbf{x})\}_{j=1}^{9}$ and $R_{\beta_3}(S_{D/L}(\mathbf{x})) = \max_{j=1,...,9} S_{D/L}^{(j)}(\mathbf{x})$ because $\beta_3 = 0.95$.

The surrogate model for lift is constructed similarly. Since $\alpha_4 = \alpha_5 = 0$, the surrogate model for lift is achieved by carrying out least-squares regression. We rely on the same data set as above, which again is divided into $J = 9$ groups that result in surrogate models $S_L^{(j)}(\mathbf{x}), j = 1,...,9$, and the overall model $S_L(\mathbf{x}) = \{S_L^{(j)}(\mathbf{x})\}_{j=1}^{9}$. Since $\beta_4 = 0$, this results in $R_{\beta_4}(S_L(\mathbf{x})) = \sum_{j=1}^{9} S_L^{(j)}(\mathbf{x})/9$. For the negative lift, we simply reverse the sign.

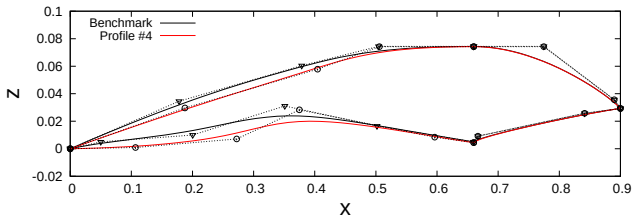

Fig. 5. Design 4 (red curve, circles for control points) compared with benchmark shape (black curve, squares for control points). Values in meters.

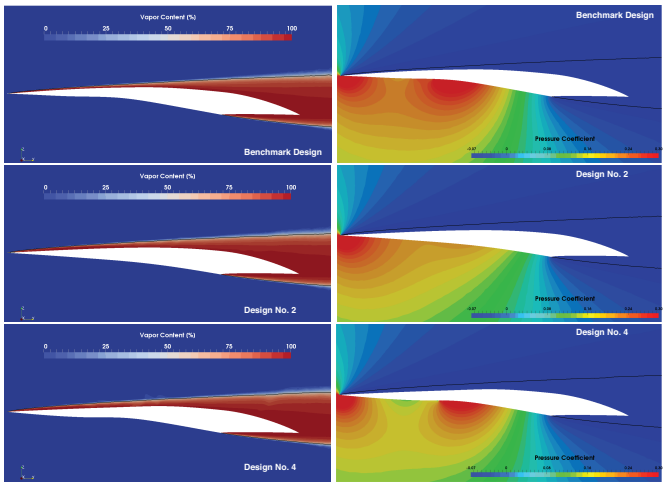

Fig. 6. Time-averaged vapor content (left column) and pressure coefficients (right column) contours. From top to bottom: benchmark design, Design 2, and Design 4. Flow conditions as in Figure 1.

5 Results and Candidate Designs

We are now in a position to generate sets of candidate designs. For the five QoIs, we have the conditions

$$
\begin{aligned}
R_{\alpha_1}(-w) &\le -w^{\min} && \text{(sufficiently large inertia)}\\
R_{\alpha_2}(-t_P^2) &\le -t_P^{\min} && \text{(sufficiently large thickness)}\\
R_{\beta_3}(S_{D/L}(\mathbf{x})) &\le 0.1 && \text{(sufficiently low drag-over-lift)}\\
R_{\beta_4}(S_L(\mathbf{x})) &\le 0.2625 && \text{(sufficiently low lift)}\\
R_{\beta_5}(-S_L(\mathbf{x})) &\le -0.23 && \text{(sufficiently high lift)}
\end{aligned}
$$

defining the candidate designs, where the first two requirements avoid surrogate models and thus R_{β_1} and R_{β_2} become immaterial. (Since $R_{\alpha_1}(-w)$ is deterministic, $R_{\beta_1}(R_{\alpha_1}(-w)) = R_{\alpha_1}(-w)$ and similarly for the second requirement.) A design $\mathbf{x}$ that satisfies these inequalities is guaranteed to meet the specified requirements in a mathematically precise sense accounting for both parametric un-

certainty in the manufacturing process of the hydrofoil and uncertainty in surrogate models.

The above five conditions are trivial to evaluate for each design $\mathbf{x}$ as they involve only explicitly or easy-to-compute quantities. No high-fidelity simulations are required. Low-fidelity output is needed to evaluate the surrogate models due to the presence of h in Eqn. (3), but they are quickly obtained. It is therefore easy to check whether a given design satisfies these conditions and also to generate a collection of candidate designs satisfying the conditions. For the sake of demonstration, we simply randomly generate 500 designs within the bounds of Table 1 and for each check the five conditions. This results in four candidate designs that satisfy the conditions. (Many more could have been obtained by further sampling and other methods.) In the following, we concentrate on two of them, labelled Design 2, with (in meters) $\mathbf{x} = (0.0834, 0.2629, 0.4258, 0.5611, 0.0040, 0.0099, 0.0147, 0.0094, 0.0027, 0.1887, 0.3929, 0.5602, 0.0300, 0.0602, 0.0742)$, and Design 4, with (in meters) $\mathbf{x} = (0.1069, 0.2724, 0.3749, 0.5962, 0.0009, 0.0072, 0.0283, 0.0084, 0.0045, 0.1877, 0.4047, 0.5058, 0.0297, 0.0578, 0.0742)$; see Figures 4 and 5. Designs 2 and 4 are quite different, with the two other candidate designs resembling these.

Table 2 summarizes the performance of Designs 2 and 4. Columns 2 and 3 give the drag-over-lift for Designs 2 and 4, respectively. Rows 3-11 give the estimates provided by the surrogate models $S^{(j)}_{D/L}(\mathbf{x})$, $j = 1, ..., 9$, resulting in the assessment $R_{\beta_3}(S_{D/L}(\mathbf{x}))$ given in the second-to-last row. The latter quantity was compared with the requirement 0.1 to determine whether the designs satisfied the requirement. For the sake of validation, we compute the actual s-risk of drag-over-lift using 898 high-fidelity simulations and obtain the results in the last row. We observe that although some individual surrogate models (for example $S^{(3)}(\mathbf{x})$) underestimate drag-over-lift, $R_{\beta_3}(S_{D/L}(\mathbf{x}))$ provides conservative estimates of the actual drag-over-lift. This highlights the importance of considering the uncertainty in surrogate models. Columns 4 and 5 give similar results for lift; again surrogate models are reasonably accurate. In comparison, the benchmark design has actual s-risk for drag-over-lift of 0.0880 and for lift of 0.2675, the latter being excessively high.

The s-risk of the inertia for the benchmark, Design 2, and Design 4 are $7.400 \cdot 10^{-6}$, $8.738 \cdot 10^{-6}$, and $8.48489 \cdot 10^{-6}$, respectively. We see that when accounting for manufacturing errors, the benchmark design fails to meet the requirement of $w^{\min} = 8.1 \cdot 10^{-6}$. The s-risk for the corresponding profile thicknesses are 0.00118, 0.00135, and 0.00185. Again, the benchmark design fails to meet the requirement of $t_P^{\min} = 0.00132$.

Figures 4 and 5 shows that the suction side control points of Designs 2 and 4 have moved to achieve a thicker cavity by lowering the back profile surface. The minimum hydrofoil thickness at 2% of the cord, required to avoid fluttering, has been recovered shifting the pressure side control points close to the leading edge of the hydrofoil. The control points of the pressure side have been changed in different ways to ensure the overall hydrodynamic performance: Design 2 has

Table 2. Performance of Designs 2 and 4.

	Drag/Lift		Lift	
Model	Design 2	Design 4	Design 2	Design 4
$S^{(1)}(\mathbf{x})$	0.0959	0.0971	0.2354	0.2367
$S^{(2)}(\mathbf{x})$	0.0937	0.0949	0.2346	0.2368
$S^{(3)}(\mathbf{x})$	0.0925	0.0930	0.2346	0.2397
$S^{(4)}(\mathbf{x})$	0.0958	0.0988	0.2341	0.2355
$S^{(5)}(\mathbf{x})$	0.0954	0.0982	0.2352	0.2368
$S^{(6)}(\mathbf{x})$	0.0954	0.0990	0.2323	0.2337
$S^{(7)}(\mathbf{x})$	0.0962	0.0989	0.2345	0.2381
$S^{(8)}(\mathbf{x})$	0.0932	0.0951	0.2345	0.2381
$S^{(9)}(\mathbf{x})$	0.0956	0.0970	0.2363	0.2390
$R_\beta(S(\mathbf{x}))$	0.0962	0.0990	0.2346	0.2371
Actual	0.0956	0.0971	0.2396	0.2465

a single curvature shape with a maximum at about $x = 0.42$ m, while Design 4 shows general higher curvatures with a marked slope change at about $x = 0.28$ m. Design 4 has pressure side shape similar to that of the benchmark design, but the maximum curvature of the concave part is shifted towards the trailing edge. Design 2 has a single curvature face and Design 4 a double curvature face. Design 2 is similar to the Johnson two-terms profile [26], also known as the Tulin-Burkat hydrofoil [27]; the single curvature pressure surface leads to lower lift force generation, which for this particular shape has been compensated through the vertical displacement of the face cavitator. Moving the sharp edge on the pressure side has the direct effect of changing the cavity detachment position and virtually increasing the operating angle of attack. Design 4 is characterized by a double curvature resembling a Johnson five-terms profile and leading to lower drag-over-lift, i.e., improved efficiency. As shown in Figure 5, the vertical position of the face cavitator does not move compared with the benchmark shape. Hence, the increase in lift generation is mainly due to the double curvature leading to a double pressure peak as evidenced in Figure 6, which shows pressure as well as vapor content for all there designs. Design 4 experiences a second pressure peak upstream the face cavitator, locally behaving as a stagnation point. The presence of bump on the face in between the two pressure peaks generates, as seen in Figure 6, a lower pressure region within the maximum curvature points. Similarly to the situation for five-term profiles, this region could be of interest in certain flow conditions. Both Designs 2 and 4 present a reduced suction surface curvature with respect to the benchmark design. This feature directly affects cavity thickness close to the leading edge as seen in Figure 6 and, hence, ensures more robust cavitating regimes than for the benchmark profile. The set-based design framework under uncertainty thus generated designs that have slightly lower

efficiency due to a larger cavity thickness, but that are more robust to changes in geometry due to manufacturing errors.

Acknowledgements

This work in supported in parts by DARPA under N66001-15-2-4055.

References

[1] Singer, D. J., Doerry, N., and Buckley, M. E., 2009. "What is set-based design?". *ASNE Naval Engineers Journal,* **4**(121), pp. 31–43.

[2] Rockafellar, R. T., and Uryasev, S., 2000. "Optimization of conditional Value-at-Risk". *Journal of Risk,* **2**, pp. 493–517.

[3] Rockafellar, R., and Uryasev, S., 2002. "Conditional value-at-risk for general loss distributions". *Journal of Banking and Finance,* **26**, pp. 1443–1471.

[4] Rockafellar, R., and Royset, J. O., 2015. "Engineering decisions under risk-averseness". *ASCE-ASME J. Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering,* **1**(2), p. 04015003.

[5] Papalambros, P. Y., 2012. "Special issue: Design under uncertainty". *ASME Journal of Mechanical Design,* **134**(10), pp. 100201–100201–1.

[6] Ayyub, B., Akpan, U., Koko, T., and Dunbar, T., 2015. "Reliability-based optimal design of steel box structures. i: Theory". *ASCE-ASME J. Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering,* **1**(3), p. 04015009.

[7] Akpan, U., Ayyub, B., Koko, T., and Rushton, P., 2015. "Development of reliability-based damage-tolerant optimal design of ship structures". *ASCE-ASME J. Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering,* **1**(4), p. 04015013.

[8] Meng, D., Li, Y.-F., Huang, H.-Z., Wang, Z., and Liu, Y., 2015. "Reliability-based multidisciplinary design optimization using subset simulation analysis and its application in the hydraulic transmission mechanism design". *ASME Journal of Mechanical Design,* **137**(5), pp. 051402–051402–9.

[9] Dutta, S., Ghosh, S., and Inamdar, M., 2016. "Reliability-based design optimization of frame-supported tensile membrane structures". *ASCE-ASME J. Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, p. G4016001.

[10] Wang, S., and Qing, X., 2016. "A mixed interval arithmetic/affine arithmetic approach for robust design optimization with interval uncertainty". *ASME Journal of Mechanical Design,* **138**(4), pp. 041403–041403–10.

[11] Royset, J. O., Der Kiureghian, A., and Polak, E., 2001. "Reliability-based optimal structural design by the decoupling approach". *J. Reliability Engineering and System Safety,* **73**(3), pp. 213–221.

[12] Royset, J. O., Der Kiureghian, A., and Polak, E., 2006. "Optimal design with probabilistic objective and constraints". *J. Engineering Mechanics,* **132**(1), pp. 107–118.

[13] Rockafellar, R. T., and Royset, J. O., 2010. "On buffered failure probability in design and optimization of structures". *Reliability Engineering & System Safety,* **95**, pp. 499–510.

[14] Brizzolara, S., and Bonfiglio, L., 2015. "Comparative cfd investigation on the performance of a new family of super-cavitating hydrofoils". In Journal of Physics: Conference Series, Vol. 656, IOP Publishing, p. 012147.

[15] Brizzolara, S., 2014. "Watercraft device". *US Patent*(US 13362298), 09.

[16] Forrester, A., Sobester, A., and Keane, A., 2008. *Engineering Design via Surrogate Modelling: A Practical Guide*. Wiley.

[17] Perdikaris, P., Venturi, D., Royset, J., and Karniadakis, G., 2015. "Multi-fidelity modeling via recursive co-kriging and gaussian markov random fields". *Royal Society Proceedings A,* **2179**(471).

[18] Royset, J. O., and Wets, R. J.-B., 2015. "Fusion of hard and soft information in nonparametric density estimation". *European J. of Operational Research,* **247**(2), pp. 532–547.

[19] Rockafellar, R. T., and Royset, J. O., 2015. "Measures of residual risk with connections to regression, risk tracking, surrogate models, and ambiguity". *SIAM J. Optimization,* **25**(2), pp. 1179–1208.

[20] Rockafellar, R. T., Uryasev, S., and Zabarankin, M., 2008. "Risk tuning with generalized linear regression". *Mathematics of Operations Research,* **33(3)**, pp. 712–729.

[21] Brizzolara, S., 2015. "A new family of dual-mode super-cavitating hydrofoils.". In Proceedings of the Fourth International Symposium on Marine Propulsors.

[22] G. Vernengo, L. Bonfiglio, S. G., and S.Brizzolara, 2016. "Physics-based design by optimization of unconventional super-cavitating hydrofoils". *Journal of Ship Research*.

[23] Hirt, C. W., and Nichols, B. D., 1981. "Volume of fluid (vof) method for the dynamics of free boundaries". *Journal of computational physics,* **39**(1), pp. 201–225.

[24] Kunz, R. F., Boger, D. A., Stinebring, D. R., Chyczewski, T. S., Lindau, J. W., Gibeling, H. J., Venkateswaran, S., and Govindan, T., 2000. "A preconditioned navier–stokes method for two-phase flows with application to cavitation prediction". *Computers & Fluids,* **29**(8), pp. 849–875.

[25] Bonfiglio, L., and Brizzolara, S., 2016. "A multiphase ranse-based computational tool for the analysis of super-cavitating hydrofoils". *Naval Engineers Journal,* **128**(1), pp. 67–85.

[26] Johnson, V. E., 1961. "Theoretical and experimental investigation of supercavitating hydrofoils operating near the free water surface". *NASA*(19980228449).

[27] Tulin, M., and Burkart, M., 1955. Linearized theory for flows about lifting foils at zero cavitation number. Tech. rep., DTIC Document.